Modeling the Interplay Between Individual Behavior and Network Distributions

Yang Yang[†], Jie Tang[†], Yuxiao Dong[‡], Qiaozhu Mei[⋆], Reid A. Johnson[‡], Nitesh V. Chawla[‡]
[†]Department of Computer Science and Technology, Tsinghua University
[‡]Department of Computer Science and Engineering, University of Notre Dame
[⋆]Department of Electrical Engineering and Computer Science, University of Michigan
yyang.thu@gmail.com, jietang@tsinghua.edu.cn, {ydong1, rjohns15, nchawla}@nd.edu, qmei@umich.edu

arXiv:1511.02562v1 [cs.SI] 9 Nov 2015

ABSTRACT

It is well-known that many networks follow a power-law degree distribution; however, the factors that influence the formation of their distributions are still unclear. How can one model the connection between individual actions and network distributions? How can one explain the formation of group phenomena and their evolutionary patterns?

In this paper, we propose a unified framework, **M^{3}D**, to model human dynamics in social networks from three perspectives: macro, meso, and micro. At the micro-level, we seek to capture the way in which an individual user decides whether to perform an action. At the meso-level, we study how group behavior develops and evolves over time, based on individual actions. At the macro-level, we try to understand how network distributions such as power-law (or heavy-tailed phenomena) can be explained by group behavior. We provide theoretical analysis for the proposed framework, and discuss the connection of our framework with existing work.

The framework offers a new, flexible way to explain the interplay between individual user actions and network distributions, and can benefit many applications. To model heavy-tailed distributions from partially observed individual actions and to predict the formation of group behaviors, we apply **M^{3}D** to three different genres of networks: Tencent Weibo, Citation, and Flickr. We also use information-burst prediction as a particular application to quantitatively evaluate the predictive power of the proposed framework. Our results on the Weibo indicate that **M^{3}D**'s prediction performance exceeds that of several alternative methods by up to 30%.

Categories and Subject Descriptors

H.3.3 [**Information Search and Retrieval**]: Text Mining; J.4 [**Social Behavioral Sciences**]: Miscellaneous

General Terms

Algorithms, Experimentation

Keywords

User behavior, Information diffusion, Heavy-tailed distributions

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...$15.00.

1. INTRODUCTION

Statistics show that 1% of Twitter users produce 50% of its content [31] and control 25% of its information diffusion [15], while 5% of Wikipedia contributors generate 80% of its content [18]. Such phenomena have received much attention, and several state-of-the-art models [3, 7, 9, 20] have been proposed, to explain their underlying mechanism. However most of these studies focus on modeling the totality of interactions between individuals, while ignoring the temporal aspects of individual actions [26]. Yet the dynamics of social phenomena are, at a fundamental level, driven by individual user actions [29], resulting in a clear and present need for understanding the connection between human dynamics and network distributions.

The connection between human dynamics (e.g., e-mail communication between individuals) and network distributions has been studied in physics [26, 29], economics [6], and sociology [4, 13]. Vázquez et al. [29] showed that the timing of individual user actions follows a non-Poisson distribution pattern, and the "bursty" nature of human behavior can be modeled based on the decisions of individual user. Rybski et al. [24] studied group behavior in social communities, and tried to understand the origin of clustering and long-term persistence. Muchnik et al. [18] tried to understand how network distributions such as degree distribution (power law) arises from individual actions. They found that action and degree are not strongly correlated. However, these studies do not provide explicit explanations for the connection between individual actions and network distributions. Recently, Song et al. [26] focused on studying communication patterns between users using mobile, e-mail, Twitter, instant message data. They discovered a series of interesting relationships that quantitatively connect human dynamics to several properties of the network.

In this work, our goal is to develop a theoretical framework to model human dynamics in social networks from three perspectives: macro, meso, and micro. At the micro-level, we try to capture how individual users make a decision to perform an action (e.g., to retweet a message on Twitter). At the meso-level, we study how individual actions develop into group behavior (e.g., the diffusion of a message) and how group behavior evolves over time. At the macro-level, we investigate how the network distributions such as power-law (or heavy-tailed phenomena) arise from group behavior.

Figure 1 illustrates the problem addressed in this paper. The left figure shows the example in which two tweets are diffused in Twitter by retweeting. In the diffusion process, each individual makes a decision to retweet or not, according to a personalized binomial distribution. The number of retweets for each message has been modeled using a lognormal distribution, and the retweet counts for all messages follow the power-law distribution. The right figure shows several potential applications, namely: modeling network

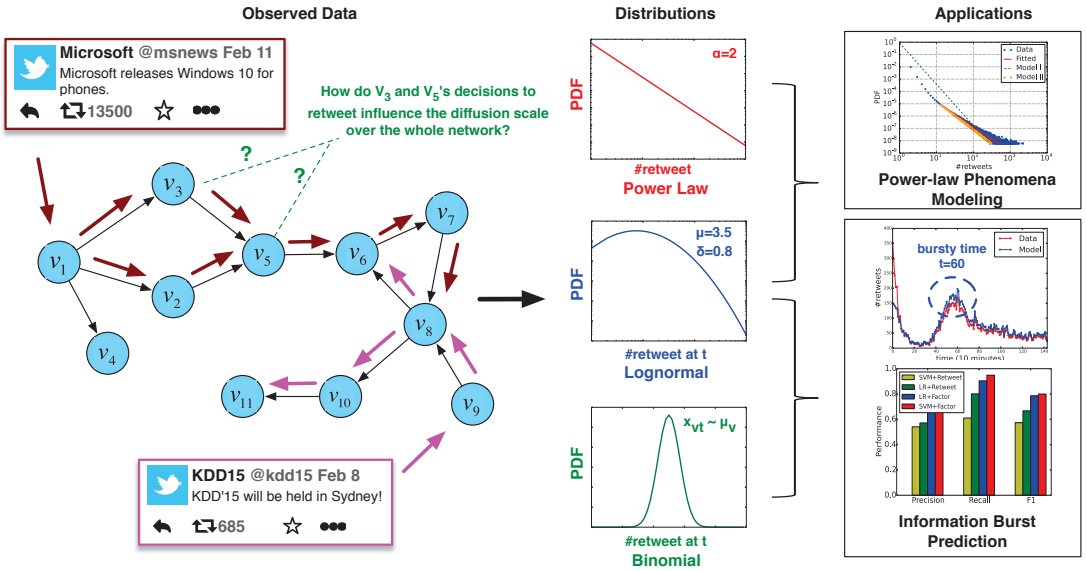

Figure 1: Example of modeling individual retweeting behavior and the distribution of the number of retweets. The left figure shows an example in which two tweets are diffused in Twitter by retweeting. The middle figure shows the binomial distribution that determines individual user actions, the lognormal distribution that capture the retweet count for each tweet, and the power law distribution that models the distribution of the number of retweets for all tweets. The right figure shows examples of applications.

phenomena using power-law and information-burst prediction. As a non-trivial problem, the fundamental challenge lies in the uncertainty of how different sub-models are intrinsically connected and how they are developed. It is also important to validate the effectiveness of such a modeling framework in real, large networks.

In this paper, we conclusively demonstrate the underlying mechanisms by which heavy-tailed phenomena develop from individual user actions. We propose a unified framework, referred to as $\mathbf{M}^3\mathbf{D}$, to 1) model the statistical distributions of individual actions, group behaviors, and heavy-tailed phenomena, and to 2) unveil the emerging process of heavy-tailed phenomena from individual actions and group behavior. This proposed framework produces several interesting results, both theoretical and empirical. Theoretically, we obtain a theorem that suggests a lower bound on the individual action-adoption probability for the existence of the power-law distribution in the number of action adopters. Empirically, by leveraging our framework, we demonstrate that it is possible to achieve an accuracy of 90% for predicting future information bursts.

The proposed framework is flexible and can benefit many applications. To model heavy-tailed distributions from partially observed individual actions, we apply $\mathbf{M}^3\mathbf{D}$ to three genres of networks: Tencent Weibo[1], Citation, and Flickr. We verify that our proposed framework can explain the emerging process of heavy-tailed phenomena from individual actions in these networks. For example, our results on the Weibo network—with more than 320 million users and 4.6 billion tweets over four months—convincingly demonstrate that 1) the retweeting action of each individual aggregates to the lognormal distribution as suggested in our framework, and 2) the lognormal distribution at each timestamp is integral to the power law distribution. Our conclusions on the connection between individual actions and heavy-tailed phenomena in real-world networks give rise to important implications for understanding the underlying mechanisms of social emergence.

[1]http://t.qq.com, one of the largest microblogging services in China.

Organization. The rest of the paper is organized as follows. In Section 2, we describe the unified framework we propose to model user actions at three levels of granularity, and describe how the three levels connect with each other. In Section 3, we introduce our experimental setup based on data from three real social network. In Section 4, we present the experimental results to validate $\mathbf{M}^3\mathbf{D}$. In Section 5, we review relevant related work. Section 6 concludes the paper.

2. MODEL FRAMEWORK

In this section, we propose a unified framework $\mathbf{M}^3\mathbf{D}$ to model the interplay between individual actions at the micro-level, group behaviors at the meso-level, and network distribution at the macro-level.

2.1 Formulation

Let $G = (V, E)$ denote an *observed* network that is a subnetwork of the *complete* network $H = (V^H, E^H)$, as H itself is too large to be observed entirely in practice. V is the set of users and $E \subset V \times V$ is the set of edges between users. Our goal is to study how individual users' actions in G emerge as the macroscopic phenomena in the complete network H. To begin with, we first give definitions of several concepts over a social network: individual actions, group behaviors, and network distribution.

Definition 1. **Individual action.** An action z performed by user v at time t is represented as $x_{tvz} \in \{-1, 1\}$. When we observe user v's action z at time t, we denote $x_{tvz} = 1$; otherwise $x_{tvz} = -1$.

The action can be defined differently in different social networks. For example, in Twitter, we can define the action as retweeting a tweet; in a scientific network, we define the action as citing a specific paper; and in Flickr, we define the action as posting comments to a specific photo. By accumulating individual actions, we can observe group behavior at the meso-level. Formally, we have the following definition.

Definition 2. **Group behavior.** For a given action z, we denote the number of users (adopters) who adopted the action at a specific time t in the complete network H as $n_z(t)$.

It is worth noting that the group behavior is defined over the complete network H instead of the observed network G. Also we use $N_z(t)$ to denote the cumulative number of adopters up to time t, and N_z to denote the total number of adopters until time $T \geq t$ (T is larger than any observed time t). Finally, at the macro-level, we consider all actions $\{z\}$ in the social network.

Definition 3. **Network distribution.** Given a set of actions $\{z\}$, and the corresponding set of numbers $\{N_z\}$ for all the actions, $P(N_z)$ represents the network distribution of the actions.

Regarding the network distribution, heavy-tailed distributions have been demonstrated to be ubiquitous in social networks [20, 7, 9, 3]. In this work, we focus on modeling the *heavy-tailed phenomena* as a macro-level reflection of individual actions in social networks. Finally, as a conclusion, an individual action is assigned to each user in the observed network G and represents the state of the user (as having adopted the action or not). A group behavior is assigned to a group of users and represented as the *number* of action adopters within the group. A heavy-tailed phenomenon is represented as a *distribution* to describe the popularity of each action over the complete network H.

Goal. Given an observed network G that is a subnetwork of the complete network H, how to unveil the mechanisms by which the heavy-tailed phenomena in network H emerge from individual actions x_{tvz} in network G? We propose a unified framework to model 1) individual actions, group behaviors, and heavy-tailed phenomena together; and 2) the emerging process by which individual actions x_{tvz} of users are revealed to be integral to the heavy-tailed phenomena in network H as a whole.

2.2 Modeling Individual Actions and Group Behavior

We first introduce independent models to model individual actions and group behavior. We then demonstrate how group behavior arises from individual actions.

For each user v in the observed network G, and a given action z, it is natural to assume individual actions x_{tvz} at different timestamps t are independent and identically distributed (i.i.d.). Given this, we formally define the individual model of user v as follows:

Definition 4. **Individual Model**. In the observed network G, each user v is assigned a binomial parameter μ_{vz} for each action z, such that for any time t, we have $x_{tvz} \sim Bernoulli(\mu_{vz})$ with $P(x_{tvz} = 1) = \mu_{vz}$, $P(x_{tvz} = -1) = 1 - \mu_{vz}$.

We then define the following group model to connect individual actions and group behavior.

Definition 5. **Group Model.** At time t, let y_t^+ be the number of individual actions with $x_{tvz} = 1$ and y_t^- be the number of individual actions with $x_{tvz} = -1$. The group model is then a dynamic system governing the evolution of group behavior over time:

$$n_z(t) = n_z(t-1) U_z^{y_t^+} D_z^{-y_t^-}, \tag{1}$$

where U_z is an "upward factor", which describes how an individual user in G adopting z will influence others in H; D_z is a "downward factor", which describes how an individual user who did not adopt z influences the others not to adopt the action.

Additionally, $n_z(t_0) = 1$, where t_0 is the timestamp when the first user adopts z. Without further explanation, we refer to t_0 as timestamp 0, for the sake of simplicity. Please notice that both the individual model and the group model are defined based on the action z. We omit the subscript z in μ_{vz}, U_z and D_z in the following descriptions to keep our notations simple.

Behavior of M^{3}D: a lognormal arises. Let us now examine our framework to see how group behavior distributes under a certain configuration of a group model's parameters.

Theorem 1. With the definition of $U = D = e$, when t is large enough, $n_z(t)$ converges to a lognormal random variable with a mean $t\sum_v \mu_v$ and a variance $t\sum_v \mu_v(1-\mu_v)$.

PROOF. We first consider another representation of y_t as

$$y_t = y_t^+ - y_t^-. \tag{2}$$

The total number of users m can be denoted as $m = y_t^+ + y_t^-$. Together with Eq. 2, we have

$$y_t^+ = \frac{1}{2}(m + y_t), \quad y_t^- = \frac{1}{2}(m - y_t). \tag{3}$$

With the definition $U = D = e$, according to Eq. 1, we have

$$\begin{aligned} n_z(t) &= n_z(t-1)\exp(\frac{m+y_t}{2})\exp(\frac{-m+y_t}{2}) \\ &= n_z(0)\exp(\sum_t y_t). \end{aligned} \tag{4}$$

This is equivalent to proof that when t is large enough, $\sum_t y_t \to N(-t\times\sum_v \mu_v, -t\times\sum_v(\mu_v)(1-\mu_v))$. We use $y'_v = \sum_{t'=1}^t x_{vt'}$ to denote the total effects made by user v. As $\{x_{vt'}\}$ are i.i.d, when t is large enough, according to the central limit theorem, we have

$$\sqrt{t}[\frac{1}{t}\sum_{t'=1}^{t} x_{vt'} - \mu_v] \sim N(0, \mu_v(1-\mu_v)). \tag{5}$$

Thus we obtain

$$y'_v = \sum_{t'=1}^{t} x_{vt'} \xrightarrow{d} N(t\mu_v, t\mu_v(1-\mu_v)) \tag{6}$$

and

$$\sum_t y_t = \sum_{v=1} y'_v \xrightarrow{d} N(t\sum_v \mu_v, t\sum_v \mu_v(1-\mu_v)). \tag{7}$$

Therefore, $n_z(t) \propto \exp(y_t)$ converges to the lognormal distribution with mean $t\sum_v \mu_v$ and variance $t\sum_v \mu_v(1-\mu_v)$, i.e.,

$$P(n_z(t)) = \frac{1}{n_z(t)\delta\sqrt{2\pi t}}\exp(\frac{-(lnn_z(t) - \tau t)^2}{2\delta^2 t}), \tag{8}$$

where

$$\tau = \sum_v \mu_v, \quad \delta^2 = \sum_v \mu_v(1-\mu_v). \tag{9}$$

□

We refer to τ and δ as the parameters of the group model. Together with time t, they express the lognormal distribution that arises for n_z. More generally, approximate lognormal distributions can be obtained when $U = D = e$ does not holding. Specifically, consider

$$\ln n_z(t) = \ln n_z(0) + \sum_{t'=1}^{t} \ln U^{y_{t'}^+} + \sum_{t'=1}^{t} \ln D^{y_{t'}^-}. \quad (10)$$

According to the Central Limit Theorem, $\sum_{t'=1}^{t} \ln U^{y_{t'}^+}$ and $\sum_{t'=1}^{t} \ln D^{y_{t'}^-}$ converge to a normal distribution. As the product of lognormal distributions is again lognormal, for sufficiently large t, $n_z(t)$ will asymptotically approach a lognormal distribution.

Significance. We conclusively demonstrate that the group behavior of the network—the collection of random (binomial) individual actions—follows a lognormal distribution. Specifically, the parameters of the lognormal distribution can be represented by our individual models, i.e., the lognormal distribution at time t is parameterized with the mean as $t\sum_v \mu_v$ and the variance as $t\sum_v \mu_v(1-\mu_v)$, where μ_v is the binomial parameter from the individual model.

2.3 Modeling Heavy-Tailed Phenomena

We study how the integration of user behavior over time eventually exhibits the heavy-tailed phenomena in the complete network.

For each action z, we define N_z as the total number of adopters in the complete social network H. Formally, let T denote the observation time window; we have $N_z = \int_{t=0}^{T} n_z(t)$. We then study the distribution of N_z.

As we concluded above, $n_z(t)$ converges to a lognormal variable when t is sufficiently large, i.e., $\ln P(n_z(t)) \sim N(\phi, \varphi)$, where ϕ and φ denote the mean and the variance, respectively. We assume $U = D = e$ holds. Therefore, $\phi = \tau t$ and $\varphi = \delta^2 t$ ($\tau = \sum_v \mu_v$, $\delta^2 = \sum_v \mu_v(1-\mu_v)$). Further assuming the observation time window T is weighted exponentially with parameter λ, we have the following theorem.

Theorem 2. $P(N_z) = CN_z^{\alpha}$, where $C = \frac{\lambda}{\delta\sqrt{(\frac{\tau}{\delta})^2+2\lambda}}$ and $\alpha = -1 + \frac{\tau}{\delta^2} - \frac{\sqrt{(\tau^2+2\lambda\delta^2)}}{\delta^2}$.

PROOF. For action z, $P(N_z)$ is equivalent to the mixture of group models whose observation time parameter T is weighted exponentially. Formally, we have

$$P(N_z) = \int_0^{\infty} \lambda \exp(\lambda t) \frac{1}{N_z\delta\sqrt{2\pi t}} \exp(\frac{-(lnN_z - \tau t)^2}{2\delta^2 t})dt. \quad (11)$$

Let $a = \lambda - \frac{\tau^2}{2\delta^2}$, $b = -\frac{(\ln n_d)}{2\delta^2}$, and $c = \frac{\lambda}{n_d\sqrt{2\pi\delta^2}\exp(\frac{\tau \ln n_d}{\delta^2})}$. It is obvious that $a, b \leq 0$. Then we have

$$\begin{aligned} Eq.\ 11 &= c\int_0^{\infty} \frac{1}{\sqrt{t}} \exp(at + \frac{b}{t})dt \\ &= 2c\int_0^{\infty} \exp(at^2 + \frac{b}{t^2})dt^2. \end{aligned} \quad (12)$$

Let $s = t^2\sqrt{\frac{b}{a}\cdot\frac{1}{t^2}}$, we have

$$Eq.\ 12 = c\exp(2\sqrt{ab})\int_{2(\frac{a}{b})^{\frac{1}{4}}}^{\infty} \frac{\exp(as^2)}{\sqrt{s^2 - 4\sqrt{\frac{b}{a}}}} ds^2. \quad (13)$$

Let $u = s^2 - 4\frac{b}{a}$, we have

$$\begin{aligned} Eq.\ 13 &= c\exp(2\sqrt{ab})\int_0^{\infty} \frac{\exp(au - 4\sqrt{ab})}{\sqrt{u}} du \\ &= \frac{c\exp(-2\sqrt{ab})}{\sqrt{-a}}\Gamma(\frac{1}{2}) \\ &= \frac{\lambda}{N_z\sqrt{2\delta^2}\sqrt{\frac{\tau^2}{2\delta^2} - \lambda}} \exp(\frac{\tau}{\delta^2}\ln N_z - 2\frac{\ln N_z}{\delta}\sqrt{\frac{1}{2}(\frac{\tau^2}{2\delta^2} - \lambda)}) \\ &= \frac{\lambda}{\sqrt{(\tau^2 - 2\delta^2\lambda)}} N_z^{-1+\frac{\tau}{\delta^2} - \frac{\tau^2 - 2\lambda\delta^2}{\delta^2}} \end{aligned} \quad (14)$$

Hence, we have $P(N_z) = CN_z^{\alpha}$, where

$$\begin{aligned} C &= \frac{\lambda}{\sqrt{(\tau^2 - 2\delta^2\lambda)}}, \\ \alpha &= -1 + \frac{\tau}{\delta^2} - \frac{\tau^2 - 2\lambda\delta^2}{\delta^2}. \end{aligned} \quad (15)$$

□

For $\alpha < 0$, we say N_z is power-law distributed. It is worth noting that, when group behavior follows a lognormal distribution, even without the conditions of $\phi = \tau t$ or $\varphi = \delta^2 t$, the power-law result still holds. The proof can be obtained by extending the proof for Theorem 2. A similar study was conducted in [1].

Behavior of M^{3}D: when does the "winner take all"? Let us now examine the behavior of our framework to see when the winner-take-all mechanism holds and leads to the power-law phenomena. Consider a simple system, in which each user shares the same individual model with the parameter p to perform an action. It turns out there is a lower bound on p for the existence of the power-law distribution over N_z.

Theorem 3. N_z is power-law distributed when $p > \frac{1}{m}$, where m is the number of users in G.

PROOF. According to Eq. 9, we have $\tau = mp$ and $\delta = mp(1-p)$. Also, according to Eq. 15, we obtain that $\beta = -1 + \frac{1-mp}{1-p} + 2\lambda$. The power-law holds for N_z when $\alpha < 0$, that is,

$$\begin{aligned} \alpha < 0 &\Rightarrow -1 + \frac{1-mp}{1-p} + 2\lambda < 0 \\ &\Rightarrow p > \frac{2\lambda + 2}{m+1} > \frac{1}{m}. \end{aligned} \quad (16)$$

□

When does the "winner take all"? Assuming we have a purely random ($p = 0.5$) system for the evolution of group behavior, as m is usually large, we are safe to claim that a power-law holds. However, considering a deterministic system, in which no user will adopt any action ($p \to 0$), the power-law will fail.

Through the above discussions, we are given some insight into the fact that randomness and the aggregate effect of individual actions finally result in the macroscopic-power-laws phenomena.

Significance. We provide evidence of how heavy-tailed phenomena in social networks emerge from individual actions. Specifically, a power-law distribution is determined by the parameters derived

Table 1: Statistics of the datasets.

	Weibo	Citation	Flickr
Users	326,497,021	1,712,431	259,565
Posts/Papers/Photos	4,634,168,136	2,092,356	854,734
Users' relations	3,274,895,719	6,485,521	1,898,069
Action logs	1,026,243,542	8,012,227	3,884,739
Time period	4 months	79 years	12 months

from our individual models and group models. We provide theoretical conditions under which a power-law distribution can form in social networks.

2.4 Further Discussions

Application. We discuss some potential applications of the proposed framework. One can integrate any machine learning algorithm into the individual model. For instance, when studying the information diffusion process in Twitter, for modeling user v's retweeting behavior, we can define features (e.g., the likelihood of the user's interests matching with the tweet's topics, profiles of the user, etc.) and construct a feature vector $\mathbf{s_{vz}}$. We then use a Logistic function $f(\cdot)$ with s_{vz} to represent the individual model μ_{vz}. With efficient training samples, we are allowed to estimate $f(\cdot)$'s parameters using Maximum Likelihood Estimation (MLE) [2]. Notice that the logistic function can be replaced by any other classification or regression models. After that, we are able to generate the "adopt" decisions of individual users, for the instantiation of the group model according to Eq. 1, and demonstrate how the retweets each tweet receives evolve over time. At last, we can calculate the parameters of the heavy-tailed distributions to present the macroscopic phenomena.

Connection with previous work. The proposed framework can be viewed as a generalization of several existing models. In Eq 1, when $U = D$, the group model can be viewed as a generalized Black-Scholes option pricing model [6]. The connection with individual models and group models is a natural multiplicative process [16]. When group behaviors follow lognormal distributions, while the conditions of $\phi = \tau t$ or $\varphi = \delta^2 t$ are not satisfied, the integration of group behaviors with heavy-tailed phenomena is similar to that of the evolutionary process of sites on the Web [1].

3. EXPERIMENTAL SETUP

3.1 Datasets

We verify the proposed framework on three different genres of large datasets: Tencent Weibo, Citation, and Flickr, Statistics of the three datasets are summarized in Table 1.

Tencent Weibo [34]. It is one of the most popular microblogging services in China. The dataset consists of 326,497,021 users, 3,274,895,719 following relationships, and 4,634,168,136 tweets, spanning over 4 months between Oct. 1st, 2011 and Jan. 30th, 2012. In this dataset, we define the user's retweeting behavior as the individual action.

Citation [28]. It is from ArnetMiner[2].This dataset consists of 2,092,356 papers published during 1950 and 2012. From those papers, we derive 1,712,433 authors, and 8,012,227 citation relationships between papers. In this dataset, the individual action is defined by a paper's citation behavior (determined by the authors in the course of scientific writing).

Flickr [33]. It was crawled from Flickr. This dataset contains 854,734 photos and 3,884,739 comments generated by 259,565 users. The dataset is used to investigate user commenting actions in photo sharing networks. Specifically, we define the individual action as whether a user posts a new comment to a specific photo.

3.2 Data Analysis

Group behavior. We examine how group behavior $n_z(t)$, such as #retweets, #citations, and #comments at each timestamp t, on real datasets distributes.

To test the hypothesis that the group behavior at a particular timestamp is lognormally distributed, we utilize QQ-plot [30], which is commonly used to compare two distributions by plotting their quantiles against each other. Specifically, for timestamp t, we first estimate the parameters of empirical normal distribution of $\ln n_z(t)$, by a MLE method, which is also used in [35]. We then plot the quantiles of $\ln n_z(t)$ in real data against certain quantiles of the estimated normal distribution.

Figure 2 shows the results on all datasets. For instance, in Figure 2(a), we test the distribution form of #retweets that a tweet receives at the 3rd hour since that tweet is posted in Weibo network. The approximate linearity of the plotted points suggests #retweets, at the 3rd hour since the original tweet is posted, follows a lognormal distribution. Analogously, we observe similar results on #citations and #comments in the Citation and Flickr networks, as shown in Figures 2(b) and (c). We also archieve similar results at other timestamps on all datasets.

Heavy-Tailed distribution. We now examine the network distribution, $P(N_z)$, on real datasets. To do so, in Weibo (Citation or Flickr) network we plot #retweets (#citations or #comments) each tweet (paper or photo) receives within 4 months (79 years or 12 months) in Figure 3. The linearity on log-log scales suggests a heavy-tailed distribution on all datasets.

We try to fit power-law distributions on the three datasets by a classical fitting technique [7] which determines two parameters: a truncation point x_{min} that governs the lower bound, above which the observed data obey a power-law distribution, and the exponent parameter α of the potential power-law distribution.

We present the power-law distributions that best fit the data in each network in Figure 3 by blue lines. We observe that the power-law distributions in the Weibo network with exponent 2.0, the Citation network with exponent 2.37, and the Flickr network with exponent 1.4. The corresponding truncation points x_{min} in each network are 13, 20, and 4, respectively. We use the p-value [7] and Residual Sum of Squares (RSS) [8] together to test the power-law hypothesis. We find that the p-values are all greater than 0.1. The average value of RSS on each dataset is 2.3×10^{-6} (see details in Table 2). The p-value, together with the small valued RSS scores, suggest that the real-world data exhibits power-law distributions [7]. We also try to fit lognormal distributions to the observed data. However, we achieve negative results on all datasets.

3.3 Evaluation Measures

It is difficult to find a ground truth to evaluate the accuracy of the proposed framework. As the most important feature of $\mathbf{M^3D}$ is the connection of macro-level network distribution with meso-level group behavior and micro-level individual actions; for quantitatively evaluation, we apply $\mathbf{M^3D}$ to fit the (macro-level) heavy-tailed distribution from observed individual actions. One advan-

[2]http://aminer.org/citation

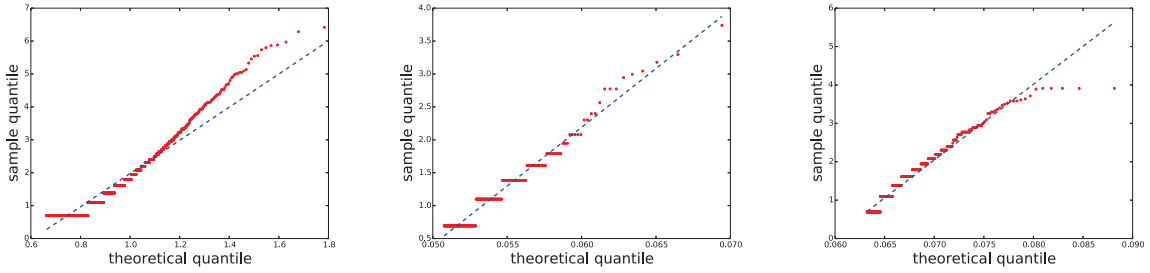

(a) 3rd hour on Weibo (b) 9th year on Citation (c) 23rd hour on Flickr

Figure 2: QQ-plots for analyzing the distribution of group behavior at a particular timestamp on each dataset.

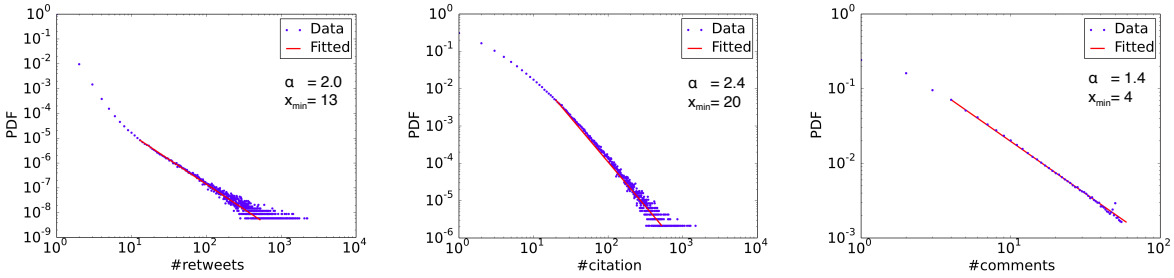

(a) (b) (c)

Figure 3: Heavy-Tailed Phenomena in Real-World Networks. Blue dots are observed data. Red lines are power-law distributions which best fit the data, with estimated exponent parameters and truncation points represented in the figures. All figures are plotted on a log-log scale.

tage of $\mathbf{M^3D}$ is that it can fit a network distribution from partially observed data. The other advantage is that we can use the group behavior or individual actions to explain the formation of the fitted distributions.

Another idea to evaluate $\mathbf{M^3D}$ is to apply it to some prediction task. $\mathbf{M^3D}$ can better capture group behavior in social networks. Thus we apply $\mathbf{M^3D}$ to information burst prediction and evaluate the prediction performance in terms of Precision, Recall, and F1-Measure.

4. EXPERIMENTAL RESULTS

To quantitatively validate $\mathbf{M^3D}$, we consider the following three evaluation aspects:

- **Fitting heavy-tailed phenomena using partially observed actions:** We thus examine to what extent $\mathbf{M^3D}$ can capture the emerging process of heavy-tailed phenomena by using partially observed individual actions in real social networks.
- **Group behavior prediction:** By this we examine the extent to which $\mathbf{M^3D}$ can model the aggregate effect of group behaviors from individual actions in real social networks.
- **Information burst prediction:** Finally, we use this application to further demonstrate the effectiveness of $\mathbf{M^3D}$.

4.1 Fitting Heavy-Tailed Phenomena using Partially Observed Actions

This task is to demonstrate whether $\mathbf{M^3D}$ can capture the emerging process of heavy-tailed phenomena by using only partially observed individual actions in real social networks.

Table 2: Residual Sum of Squares (RSS) for fitting heavy-tailed phenomena.

	Weibo	Citation	Flickr
Fitted	6.180×10^{-13}	1.426×10^{-6}	5.606×10^{-6}
Model I	8.240×10^{-1}	7.887×10^{-2}	3.867×10^{-2}
Model II	1.268×10^{-10}	7.366×10^{-7}	7.000×10^{-4}
From Group	8.062×10^{-13}	1.721×10^{-6}	1.059×10^{-5}

Problem. Given an observed subnetwork $G = (V, E)$ of a complete network H over a time window T and a set of individual actions $X = \{x_{tvz} | t \leq T, v \in V\}$ in G, the goal is to estimate whether N_z in real networks—the number of users who perform the action z in the complete network H over the observed timespan T—follows the power-law distribution parameterized by our framework.

Specifically, in the Weibo network, $x_{tvz} = 1$ indicates that user v retweets a particular tweet, indicated by z, at timestamp t. In Citation, it means author v cites another a particular paper z at time t. Analogously, the action that user v posts a comment to photo z at time t is denoted by $x_{tvz} = 1$ in the Flickr network.

Setup. In Section 3, we conclude that all three networks exhibit power-law distributions and also provide the estimated exponent parameters that best fit the real data. We use Residual Sum of Squares (RSS) [8] to quantify the distance between the distributions of N_z in real data and the distributions provided by our framework. A smaller RSS represents a better-fitted distribution.

We introduce how we generate the observed network G in different datasets. In the Weibo network, we choose the users who retweet a tweet $z \in Z$, where Z is the set of all tweets exposed

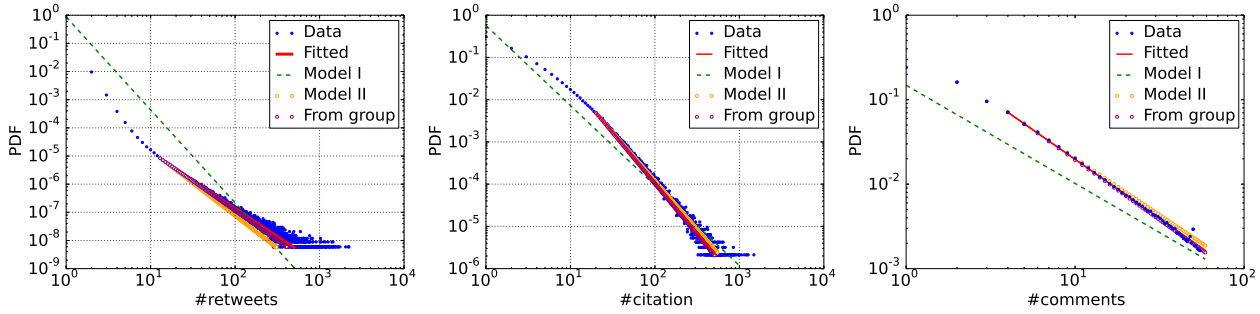

Figure 4: Fitting the heavy-tailed phenomena in (a) Weibo, (b) Citation, and (c) Flickr. All figures are plotted on a log-log scale. In each figure, x-axis denotes the number of retweets (citations or comments). Y-axis denotes the fraction of tweets (papers or photos). Blue solid dots are observed data. The red solid line is the result of fitting the observed data. The green dashed line is the result obtained by our framework, using partially observed individual actions. Yellow hollow dots are obtained by our framework considering truncation points. Purple hollow dots represent the fitting results using group behaviors.

to users V in G, within the first 50 minutes since z was posted, and the followers/followees of these users. In Citation, G consists of authors who cite papers within the first year since these papers are published. In Flickr, the users who posted comments to photos within the first hour since the photos are posted are chosen as the subnetwork G.

We introduce two methods to apply our framework to estimate the exponent parameter α of power-law distributions from individual actions in real data.

Model I. We model individual actions by a binomial model with parameter μ_{vz} in Definition 4. There are different ways to represent and estimate μ_{vz}, such as using a logistic regression to represent μ_{vz} and estimating the regression parameters by the Maximum Likelihood Estimation (MLE) method according to the individual action logs of user v. In this work, to keep our framework flexible and general, we use a straightforward method to represent μ_{vz} — that is, the probability that user v is influenced by one of her neighbors to perform the action z. Specifically, we define μ_{vz} as

$$\mu_v = \frac{\sum_t x_{tvz}}{\sum_t \sum_{u|e_{vu}\in E} x_{tuz}} \tag{17}$$

Given mu_{vz}, we then calculate τ and δ—the parameters of group behaviors in Eq. 9. Unfortunately, there is no effective way to estimate λ (the exponential parameter of the observation time window T) automatically. Thus we define λ manually and leave the automatic estimation method to our future work. In practice, we empirically set λ as 0.008, 0.1, and 0.23 in Weibo, Citation, and Flickr, respectively. We finally obtain the exponent α according to Eq. 15.

Model II. As prior work [7] and Figure 3 suggest, however, few empirical phenomena in practice obey power-laws for all observed data x. More often the power-law applies for values greater than some minimum point $x_{\min}$, which can be understood as the truncation point of the empirical phenomena. We first calculate the truncation point $x_{\min}$ by the MLE method described in [7]. We then use the posts whose retweeting numbers are no less than $x_{\min}$ to estimate parameters τ and δ by following the steps in Model I. After that, when normalizing the estimated PDF, we also only consider the tweets whose retweets are no less than the lower bound.

Results. Figure 4 shows the results on three networks. All figures are plotted on log-log scales. Blue dots represent the distribution of real data. The red solid line is the power-law fitting by observing blue dots through the MLE method [7]. Together with the real data, the fitted result can be used as ground truth. The results of Model I and Model II provided by our framework are represented by the green dashed line and yellow hollow dots, respectively.

From Figure 4, the power-law distribution suggested by our framework (Model II) is clearly seen to be a good fit to the real data (blue points). Please recall that our framework only observes parts of the complete networks (3.4% in Weibo, 1.3% in Citation, and 4.7% in Flickr). We also validate the results by quantifying the difference of distributions between Model II and real data via RSS using geometric binning. Table 2 reports the RSS results of corresponding methods. We can see that the performance of our framework can achieve RSS values at the orders of magnitude 10^{-10}, 10^{-7}, a nd 10^{-4} in the three networks, which indicates excellent numerical performance of the power-law fitting. We also notice that the estimated truncation points help our framework (Model II) better fit the heavy-tailed phenomena than Model I.

We conclude that the emerging process of heavy-tailed phenomena can be modeled and explained well from the partially observed individual actions by our framework, on all datasets.

Fitting Heavy-Tailed Phenomena using Group Behavior Next, we validate whether our proposed $\mathbf{M}^3\mathbf{D}$ is able to capture the emerging process of heavy-tailed phenomenas from group behavior. Thus we conduct another heavy-tailed phenomena fitting task: Given a set of actions $\{z\}$ and a group behavior $n_z(t)$ for each action z at each timestamp t, the goal is to estimate the network distribution N_z that is defined in Definition 3.

Following the theoretical results in Section 2 and empirical results in Section 3, $n_z(t)$ converges to a lognormal variable when t is sufficiently large. Moreover, $\mathbf{M}^3\mathbf{D}$ suggests that N_z behaves as power-law when $n_z(t)$ follows lognormal distributions. Thus, our general idea here is to first estimate each timestamp's corresponding lognormal distribution over $n_z(t)$. Then we fit the heavy-tailed phenomena, based on the estimated lognormal. Specifically, for each timestamp t, we estimate the corresponding lognormal parameters by following the method introduced in [35]. We then calculate the exponent parameter α according to Theorem 2. Please notice that we keep λ, the weighted parameter of the observation window, at the same setting as at the last task when calculating α.

We demonstrate the fitting results in Figure 4, by purple lines. We also present the RSS scores in Table 2. As we can see, compar-

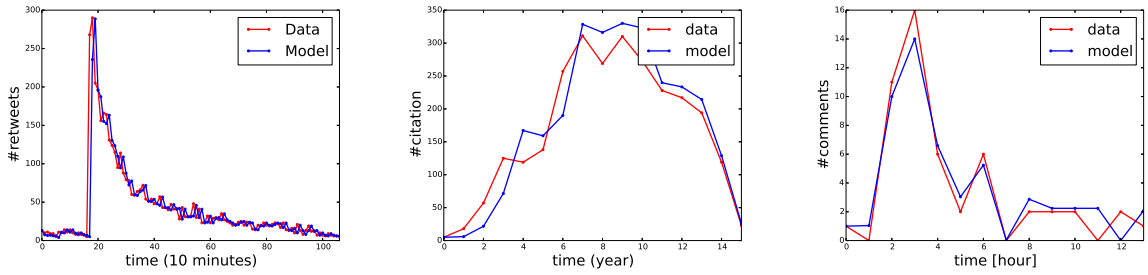

(a) (b) (c)

Figure 5: Group behavior prediction. In each figure, the x-axis denotes the time that has passed since the source tweet (paper or photo) was published. The y-axis denotes the number of retweets (citations or comments). We demonstrate the observed data (truth) and the prediction results (model).

ing with the individual actions (Model II), group behaviors provide more precise information about the lognormal parameters and obtain a better modeling result, which suggests $\mathbf{M}^3\mathbf{D}$ bridges heavy-tailed phenomena and group behavior precisely.

4.2 Group Behavior Prediction

This task is to demonstrate whether our framework can model the aggregate effect of group behaviors from individual actions in real social networks.

Problem. Given an observed subnetwork $G = (V, E)$ of a complete network H, a set of individual actions $X_t = \{x_{tvz} | v \in V\}$ at timestamp t in G, and $n_z(t)$ in network H, the goal is to infer the group behaviors $n_z(t+1)$ at the next timestamp $t+1$.

Setup. We separate the observed network G introduced in Section 4.1 into a training set and a test set. In Weibo, for each tweet z, we regard $\{n_z(t) | t \leq 500 \text{ minutes}\}$ as the training set, and the balance as the test set. In Citation, for each paper z, we regard $\{n_z(t) | t \leq 5 \text{ years}\}$ as the training set, and the rest as the test set. In Flickr, for each photo z, we use $\{n_z(t) | t \leq 2\text{hours}\}$ as the training set. We then employ the training set to estimate the "upward factor" U and "downward factor" D of the group model in our proposed framework (see details in Eq. 1). We finally calculate $n(t+1)$ by using $n_z(t)$, U, and D according to Eq. 1.

There are several methods for estimating factors U and D. Here, we assume that the factors are constant over both users and time. Under this assumption, given any two timestamps t_1 and t_2, we are able to estimate the factors according to equations below:

$$\begin{aligned} \ln \widehat{U}_{t_1 t_2} &= \frac{(\ln n(t_1+1) - \ln n(t_1)) y^-_{t_2} + (\ln n(t_2) - \ln n(t_2+1)) y^-_{t_1}}{y^+_{t_1} y^-_{t_2} - y^+_{t_2} y^-_{t_1}} \\ \ln \widehat{D}_{t_1 t_2} &= \frac{(\ln n(t_1+1) - \ln n(t_1)) y^+_{t_2} - (\ln n(t_2) - \ln n(t_2+1)) y^+_{t_1}}{y^-_{t_1} y^+_{t_2} - y^-_{t_2} y^+_{t_1}} \end{aligned} \quad (18)$$

Please note that for different pairs of timestamps, the estimated results might be different. Hence, we use the average value of all possible configurations as the final reported results. Formally, let $\widehat{U}_{t_1 t_2}$ and $\widehat{D}_{t_1 t_2}$ be the factors estimated according to time t_1 and t_2 by Eq. 18, we define $\widehat{U} = \frac{\sum_{t_1,t_2} \widehat{U}_{t_1 t_2}}{C_T^2}$ and $\widehat{D} = \frac{\sum_{t_1,t_2} \widehat{D}_{t_1 t_2}}{C_T^2}$, where T is the last timestamp in the training set.

In practice, an alternative method is to assign each user different configurations of these two factors. However, our main goal in this paper is to provide the underlying mechanisms of the emergence of heavy-tailed phenomena. Thus we keep the assumption that the two factors are independent of both time and users, to simplify and generalize our proposed framework. We leave the user- or time-dependent factor definition and estimation for our future work.

Results. We report the experimental results of group behavior prediction in Figure 5 in the three networks. In each figure, the x-axis denotes the duration since the source tweet (paper or photo) is posted by setting ten minutes (one year or one day) as the interval in Weibo (Citation or Flickr). The y-axis denotes the number of retweets (citations or comments) the source tweet (paper or photo) receives at each time interval. We plot the truth in real networks using red lines and the predicting results of $\mathbf{M}^3\mathbf{D}$ with blue lines.

Figure 5(a) presents the results for tweets with different popularity (indicated by the number of retweets) in Weibo. Clearly, we can see that the modeling results (blue lines) are well coupled with the real data (red lines) in different cases. Our method successfully captures the upward and downward tendencies of each tweet's retweeting dynamics over time precisely.

Figure 5(b) shows the results of two papers with different levels of citation counts (2262 and 122). Figure 5(c) shows the results of two photos with different numbers of comments (50 and 133). As in Weibo, group behaviors can be successfully inferred from individual citing and commenting actions in Citation and Flickr.

We conclude that our framework can capture the aggregate effect of group behaviors from individual actions, and predict the trends of dynamic popularities in real social networks.

4.3 Information Burst Prediction

We now describe ways in which to apply our framework to social applications. In this work, we focus on information burst prediction [14, 19]. Please note that the focus of the study is to demonstrate how our framework can help social applications.

Problem. Given a tweet z at timestamp t and the number of users who retweet z within the time window $[t-k, t]$ ($\{n_z(t') | t-k \leq t' \leq t\}$), the goal is to predict whether there will be an information burst at timestamp $t+1$. Formally, we say a burst happens at time t_1 if $n_z(t_1)$ is the largest in the period ranging from one hour before to one hour after t_1, i.e., $\forall t_2 \in [t_1 - 1hr, t_1 + 1\text{hr}]$, we have $n_z(t_1) > n_z(t_2)$.

Observation. In practice, the upward factor U and downward factor D, instead of being constant, may change over time. We study the trends of these factors and present the results in Figure 6. Due to space limitations, we use the factors corresponding to one tweet as an example. We observe similar results on other tweets.

In Figure 6(a), we can see that the downward factors are rel-

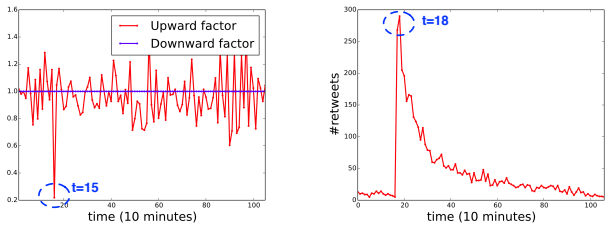

(a) (b)

Figure 6: Factor analysis. (a) The trend of the upward factor and the downward factor, corresponding to the tweet in (b), over time. (b) An example of information burst.

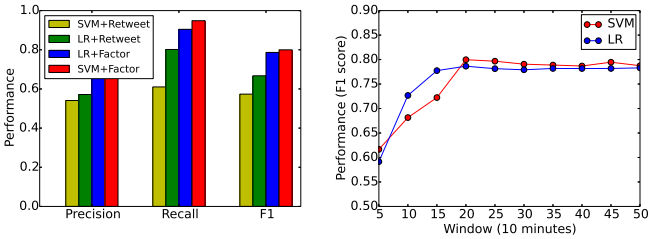

(a) (b)

Figure 7: Apply M^{3}D to information burst prediction. (a) Prediction performance by two classification models considering historic retweet times (SVM+Retweet and LR+Retweet) and models utilizing the upward factors in M^{3}D (SVM+Factor and LR+Factor). We set the observation window k as 200 minutes. (b) Prediction performance by varying the observation windows from 50 minutes to 500 minutes.

atively stable around 1.0, while the upward factors change frequently and sharply. A potential explanation is that a user's retweeting actions can influence others to retweet, while the decision not to retweet a tweet has limited effect on others' retweeting decisions. We further observe that when observing a peak or a valley in Figure 6(a), the upward factor will achieve a peak in next few timestamps in Figure 6(b). We conjecture that the burst of retweets is correlated with the variation of the corresponding upward factors.

Setup. We apply the upward factors in our framework into the information burst prediction task. We first estimate $\widehat{U}_{t'}$ at each timestamp t' in the observation time window $[t-k+1, t]$. We then use the estimated upward factors as features of classification models. We also consider a baseline method in which the number of retweets in previous steps $\{n_z(t')|t-k \leq t' \leq t\}$ is used as the features for this task. As our goal is to provide evidence of the correspondence between upward factors and information burst, we simply use classic classifiers, logistic regression (LR) and support vector machine (SVM), to report the predictability.

Results. Figure 7(b) shows the predictive performance. We can clearly see that by our methodology, both simple models significantly outperforms the comparison method (using previous retweets information—$n_z(\cdot)$ as features). Moreover, the methods that involve using upward factors derived from our framework can achieve an F1 score of 0.8, which demonstrates the predictability of bursty phenomena in social media. In terms of precision and recall, the performance is still promising.

We further examine how the length of the observation window influences the prediction performance in Figure 7(b). We observe that both LR and SVM achieve the best and stable performance when the observation window reaches 200 minutes (around three hours). We conclude that the predictability of bursty phenomena in information diffusion is highly correlated to the observation window, and information burst is more predictable when conducted over a sufficient timeframe—only three hours in Weibo.

Further study of other applications, such as cascade prediction, scientific impact modeling, and popularity forecast, is an area of work for the near future.

5. RELATED WORK

The heavy-tailed phenomena—such as power-law and lognormal distributions—have been discovered to be ubiquitous in a variety of network systems [17, 20, 7].

Power-laws have been widely observed in both nature and human society through extensive studies. Power-laws are characterized by the following probability distribution:

$$f(x) = Cx^{-\alpha}$$

where α is the exponent parameter and C is the normalization term. Essentially, power-laws model the functional relationships between two quantities, where one quantity varies as a power of another [20]. This statistical law was first revealed by the degree distributions of Internet graphs and the World Wide Web in 1999 [9, 3]. Besides having been found in Internet and WWW networks, power-law distributions have been discovered in publication citations [23], phone calls [25], tie strengths [21], and so on.

Along with power-laws, extensive studies have discovered that **lognormal** distributions are satisfied in dynamic networks. Conceptually, a lognormal distribution is defined as "a continuous probability distribution of a random variable whose logarithm is normally distributed." [17]. Formally, its density function can be expressed by the following formula:

$$f(x) = \frac{1}{x\delta\sqrt{2\pi}} \exp[-\frac{(\ln x - \mu)^2}{2\delta^2}]$$

where μ and δ are, respectively, the mean and standard deviation of the variable's natural logarithm. Huberman et al. [12] presented that the number of pages at a given site is lognormally distributed for every timestmap in the environment of WWW. Following this, similar studies also unveiled that part of the WWW pages demonstrate a truncated lognormal distribution [5, 22]. Stouffer et al. [27] discovered that the time for people to respond to emails also follows a lognormal distribution. Although a tremendous amount of exploration into the modeling of heavy-tailed phenomena in networks has been done, the mechanism of how these phenomena emerge from individual actions has received little attention.

An individual behavior, such as an action adoption, is an integral element of the heavy-tailed phenomena in social networks. A large body of studies has been focused on the modeling and predicting of individual behaviors. Xu, and Hong et al. [32, 11] modeled user online preferences and predicted individual adoption decisions in Twitter. However, the connections between individual behavior and collective dynamics are still not well studied. Recently, Rybski et al. [24] studied individual behaviors and further unveiled the origin of collective behaviors of the social community with both clustering and long-term persistence. Muchnik et al. [18] demonstrated that heavy-tailed degree distributions in networks are causally determined by similarly skewed distributions of human activity. Ghosh et al. [10] studied the interplay between a dynamic process and the structure of the network on which it is defined. The

major difference between our work and previous work lies in that we theoretically and empirically demonstrate how the integration of individual actions in social networks activates the emergence of group behavior, from which network distributions arise as a whole.

6. CONCLUSION

In this paper, we study a novel problem of modeling the interplay between individual behavior and network distributions. We propose a unified framework $\mathbf{M}^3\mathbf{D}$ to model individual behavior and network distributions together. The framework offers a way to explain how group behavior has developed and evolves over time based on individual actions, and to understand how network distributions such as power-law (or heavy-tailed phenomena) can be explained by group behavioral patterns. The framework is flexible and can benefit many applications. We apply $\mathbf{M}^3\mathbf{D}$ to three different networks: Tencent Weibo, Citation, and Flickr. Our experimental results show that $\mathbf{M}^3\mathbf{D}$ is able to model emerging network distributions in social networks from individual actions. Moreover, we use information burst prediction as an application to quantitatively evaluate the predictive power of $\mathbf{M}^3\mathbf{D}$.

7. REFERENCES

[1] L. Adamic and B. A. Huberman. The nature of markets in the world wide web. *Q. J. Econ.*, 1999.

[2] H. Akaike. Information theory and an extension of the maximum likelihood principle. In *Selected Papers of Hirotugu Akaike*, pages 199–213. 1998.

[3] A.-L. Barabási and R. Albert. Emergence of scaling in random networks. *science*, 286(5439):509–512, 1999.

[4] B. D. Bernheim. A theory of conformity. *J. Polit. Econ.*, 1027(5):841–877, 1994.

[5] Z. Bi, C. Faloutsos, and F. Korn. The "DGX" distribution for mining massive, skewed data. In *KDD '01*, pages 17–26, 2001.

[6] F. Black and M. Scholes. The pricing of options and corporate liabilities. *J. Polit. Econ.*, pages 637–654, 1973.

[7] A. Clauset, C. R. Shalizi, and M. E. Newman. Power-law distributions in empirical data. *SIAM review*, 51(4):661–703, 2009.

[8] N. R. Draper and H. Smith. *Applied regression analysis*. 2014.

[9] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the internet topology. In *COMPUT COMMUN REV*, volume 29, pages 251–262, 1999.

[10] R. Ghosh, S.-H. Teng, K. Lerman, and X. Yan. The interplay between dynamics and networks: centrality, communities, and cheeger inequality. In *KDD'14*, pages 1406–1415, 2014.

[11] L. Hong, A. S. Doumith, and B. D. Davison. Co-factorization machines: Modeling user interests and predicting individual decisions in twitter. In *WSDM '13*, pages 557–566, 2013.

[12] B. A. Huberman and L. A. Adamic. Internet: growth dynamics of the world-wide web. *Nature*, 401(6749):131–131, 1999.

[13] H. C. Kelman. Compliance, identification, and internalization: Three processes of attitude change. *J. Confl. Resolut.*, 2(1):51–60, 1958.

[14] J. Kleinberg. Bursty and hierarchical structure in streams. In *KDD '02*, pages 91–101, 2002.

[15] T. Lou and J. Tang. Mining structural hole spanners through information diffusion in social networks. In *WWW'13*, pages 825–836, 2013.

[16] M. Mitzenmacher. A brief history of generative models for power law and lognormal distributions. *Internet mathematics*, 1(2):226–251, 2004.

[17] A. M. Mood. Introduction to the theory of statistics. 1950.

[18] L. Muchnik, S. Pei, L. C. Parra, S. D. Reis, J. S. Andrade Jr, S. Havlin, and H. A. Makse. Origins of power-law degree distribution in the heterogeneity of human activity in social networks. *Scientific reports*, 3, 2013.

[19] S. A. Myers and J. Leskovec. The bursty dynamics of the twitter information network. In *WWW '14*, pages 913–924, 2014.

[20] M. E. Newman. Power laws, pareto distributions and zipf's law. *Contemporary physics*, 46(5):323–351, 2005.

[21] J.-P. Onnela, J. Saramäki, J. Hyvönen, G. Szabó, D. Lazer, K. Kaski, J. Kertész, and A.-L. Barabási. Structure and tie strengths in mobile communication networks. *PNAS*, 104(18):7332–7336, 2007.

[22] D. M. Pennock, G. W. Flake, S. Lawrence, E. J. Glover, and C. L. Giles. Winners don't take all: Characterizing the competition for links on the web. *PNAS*, 99(8):5207–5211, 2002.

[23] F. Radicchi, S. Fortunato, and C. Castellano. Universality of citation distributions: Toward an objective measure of scientific impact. *PNAS*, 2008.

[24] D. Rybski, S. V. Buldyrev, S. Havlin, F. Liljeros, and H. A. Makse. Communication activity in a social network: relation between long-term correlations and inter-event clustering. *Scientific reports*, 2, 2012.

[25] M. Seshadri, S. Machiraju, A. Sridharan, J. Bolot, C. Faloutsos, and J. Leskove. Mobile call graphs: Beyond power-law and lognormal distributions. In *KDD '08*, pages 596–604, 2008.

[26] C. Song, D. Wang, and A.-L. Barabasi. Connections between human dynamics and network science. *arXiv preprint arXiv:1209.1411*, 2012.

[27] D. B. Stouffer, R. D. Malmgren, and L. A. Amaral. Log-normal statistics in e-mail communication patterns. *arXiv preprint physics/0605027*, 2006.

[28] J. Tang, J. Zhang, L. Yao, J. Li, L. Zhang, and Z. Su. Arnetminer: extraction and mining of academic social networks. In *KDD'08*, pages 990–998, 2008.

[29] A. Vázquez, J. G. Oliveira, Z. Dezsö, K.-I. Goh, I. Kondor, and A.-L. Barabási. Modeling bursts and heavy tails in human dynamics. *Physical Review E*, 73(3):036127, 2006.

[30] M. B. Wilk and R. Gnanadesikan. Probability plotting methods for the analysis of data. *Biometrika*, 55(1):1–17, Mar. 1968.

[31] S. Wu, J. M. Hofman, W. A. Mason, and D. J. Watts. Who says what to whom on twitter. In *WWW '11*, pages 705–714, 2011.

[32] Z. Xu, Y. Zhang, Y. Wu, and Q. Yang. Modeling user posting behavior on social media. In *SIGIR '12*, pages 545–554, 2012.

[33] Y. Yang, J. Jia, S. Zhang, B. Wu, Q. Chen, J. Li, C. Xing, and J. Tang. How do your friends on social media disclose your emotions? In *AAAI'14*, 2014.

[34] Y. Yang, J. Tang, C. W.-k. Leung, Y. Sun, Q. Chen, J. Li, and Q. Yang. Rain: Social role-aware information diffusion. In *AAAI'15*, 2014.

[35] T. Zaman, E. B. Fox, E. T. Bradlow, et al. A bayesian approach for predicting the popularity of tweets. *AOAS*, 8(3):1583–1611, 2014.